

# Local Search for Optimal Global Map Generation Using Mid-Decadal Landsat Images

Robert A. Morris<sup>2</sup> and John Gasch<sup>3</sup> and Lina Khatib<sup>1,2</sup> and Steven Covington<sup>4,3</sup>

<sup>1</sup> PSGS

<sup>2</sup> Computational Sciences Division, NASA Ames Research Center

<sup>3</sup> Landsat Mission Operations, Goddard Space Flight Center

<sup>4</sup> Aerospace Corporation

## Abstract

NASA and the US Geological Survey (USGS) are generating image maps of the entire Earth using Landsat 5 Thematic Mapper (TM) and Landsat 7 (L7) Enhanced Thematic Mapper Plus (ETM+) sensor data from the period of 2004 through 2007. The map is comprised of thousands of scene locations and, for each location, there are tens of different images of varying quality to choose from. Constraints and preferences on map quality make it desirable to develop an automated solution to the map generation problem. This paper formulates a Global Map Generator problem as a *Constraint Optimization Problem* (GMG-COP) and describes an approach to solving it using local search. The paper also describes the integration of a GMG solver into a user interface for visualizing and comparing solutions.

## Introduction and Motivation

The NASA *Land-Cover and Land-Use Change (LCLUC)* Program is partnering with the USGS Earth Resources and Observation Science (EROS) Data Center to produce a high resolution mosaic map of the Earth. The map will consist of a data set of high quality images of the Earth's continental landmass using Landsat 5 (L5) Thematic Mapper (TM) and Landsat 7 (L7) Enhanced Thematic Mapper Plus (ETM+) sensor data from the mid-decadal period of 2004 through 2007. This project is known as the *Global Land Survey*, or GLS-2005.

The end-product will be composed of roughly 9500 Worldwide Reference System 2 (WRS-2) <sup>1</sup> Landsat scene locations for which there are often tens of images available to select from. Eventually, over 300,000 images must be evaluated and down-selected to create the final survey data set. The resulting data map will be distributed to the public at no charge through a USGS website. In addition to providing benefits to researchers in the Earth sciences, it

will likely become the next generation backdrop for Google-Earth (which currently uses the GeoCover-2000 data set).

A collection of diverse preference criteria defines a high quality image map. First, a good map will typically consist of the best (most cloud-free) image data available per scene. Second, each image is associated with a Normalized Difference Vegetation Index (NDVI) value, which is an historic metric of the average health and density of vegetation within that scene on the date of the image acquisition. For Earth Science applications, images with high NDVI are typically preferred. Third, to be usable for regional scientific studies, it is preferable to choose image data that are seasonally consistent with neighboring scenes. Fourth, to accommodate land-cover/land-use change analysis, consideration must be given to the seasonality of previous survey data sets. Finally, because of a malfunction in the image scanner on L7 since 2003, ETM+ produces imagery that has coverage discontinuities such that an individual image covers only 78% of the land area. To compensate, two images of the same scene taken on different days are combined to produce a composite image that partially or fully closes the gaps. Pairs of images of a common scene must therefore be chosen to maximize coverage (minimize gap), which means the two scenes should be mutually out of phase. Each image is assigned a "gap-phase statistic", or GPS, which is an absolute measure of the geometric registration of the image scan line with respect to the scene center point. Such GPS values are used to compute the area coverage of composite images.

The size of the space of global image maps, as well as the number of criteria for quality, make it desirable to automate the production of solutions to the problem. This paper presents a formulation of the map generation problem, and describes an approach to solving the problem using local search. Section 2 describes the problem in more technical detail. Section 3 describes a local search approach to solving the problem. Section 4 discusses the user interface, Landsat Scene Selector Interface (LASSI) into which the solver is integrated.

## Constraint Optimization Problem

Global Map Generation (GMG) can be viewed as a *Constraint Optimization Problem* (GMG-COP) (Larrosa & Dechter 2003), with a set of variables  $V = \{v_{i,j}\}$  indexed by WRS-2 path and row number  $i, j$ . Each variable  $v_{i,j}$

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>L5 and L7 follow the WRS-2 coordinate system for indexing locations on the Earth where data is acquired. WRS-2 indexes a location via a set of paths and rows, with a 16-day repeat cycle. L5 follows the WRS-2 system with a temporal offset of 8 days relative to L7. The WRS-2 indexes orbits (paths) and scene centers (rows) into a global grid system (daytime and night time) of 233 paths by 248 rows. We refer to each path, row element as a *scene location*.

represents a scene location, and is associated with a *domain*  $D_{i,j} = \{d_{i,j,1}, \dots, d_{i,j,m}\}$ , where each  $d_{i,j,k}$  represents a TM or ETM+ image taken of the corresponding scene. There are binary links with the 4 neighboring scenes, designated  $north(v_{i,j-1})$ ,  $east(v_{i-1,j})$ ,  $south(v_{i,j+1})$ , and  $west(v_{i+1,j})$ .

A *solution*  $s$  to the GMG COP is a set of assignments  $s = \{v_{i,j} \leftarrow \langle d_{i,j,k}, d_{i,j,l} \rangle\}$ . The need for a pair of images arises from the L7 ETM+ gap anomaly. One partial image is called the *base*; the other is called the *fill*. If the base is a TM image from L5, where there are no missing image data, we set  $d_{i,j,l} = d_{i,j,k}$  by convention. For an arbitrary solution  $s$ , we write  $b_s(v_{i,j})$ ,  $f_s(v_{i,j})$  for the base and fill values for the scene  $v_{i,j}$  assigned by  $s$ .

The GMG problem is a multiobjective optimization problem, in which a set of potentially competing preference criteria are used to evaluate and compare solutions. The preference criteria are the following:

1. Single image criteria:
  - Minimize cloud cover;
  - Maximize NDVI value;
  - Seasonality with previous data sets;
  - Relative preference for acquiring L7 versus L5 images.
  - Reward for acquisition dates centered in study period (2005 or 2006 versus the fringe years 2004 or 2007)
2. ETM+ composite criteria:
  - Minimize gaps in data that remain from compositing image pairs;
  - Minimize the difference in the days the composited images were acquired.
3. Criteria relating pairs of adjacent images:
  - Minimize the difference in the days the images were acquired;
  - Minimize the difference between the days in the year (ignoring the year) the images were acquired.
  - Prefer adjacent images acquired from a common sensor (sensor homogeneity) i.e. both TM or both ETM+.

Each candidate image is represented as a vector of “meta-data” comprised of the following text fields: the WRS-2 scene path and row numbers, the sensor that acquired the image (TM or ETM+), the acquisition date, the cloud cover assessment, the NDVI metric, and a GPS value used for evaluating the scene area coverage that results from compositing two images. In addition, a preference date is given. This date is to be used in future work to direct the solver toward solutions with minimum perturbations from these preferred dates.

Each meta-data element is associated with a function that is used to evaluate solution quality. We normalize by considering *merit* values in the range  $[0, 1]$  where 0 is worst and 1 is best. This way the objective function is a maximization and always positive. The quality of an individual image can be depicted in terms of two functions, related to the NDVI merit value, and to cloud cover:  $ndvi : D \rightarrow [0, 1]$ , and  $acca : D \rightarrow [0, 1]$ . Second, there are two functions

associated with measuring the time difference between the acquisition of neighboring pairs of images: *Absolute Day Difference*,  $date\Delta : D \times D \rightarrow [0, 1]$  is the number of days between image acquisition, and *Day of Year Difference*,  $doy\Delta : D \times D \rightarrow [0, 1]$  is the gap in days (ignoring the year in which it was acquired). The latter function is used to reward solutions that assign images that are seasonally similar, regardless of year, whereas the former rewards solutions with pairs of images taken in the same year. Third, the function *Area Coverage*,  $cover : D \times D \rightarrow [0, 1]$  assigns a value that indicates goodness of fit between a base and fill image used in a composite. Finally, to express relative preferences for TM or ETM+ images, the function  $IsL5 : D \rightarrow \{0, 1\}$ ,  $IsL7 : D \rightarrow \{0, 1\}$  assign 1 to images acquired by TM (respectively, ETM+), and 0 otherwise.

The WRS organization of scene locations into path and row induces a grid or lattice structure to the GMG-COP constraints. Because of the symmetry of adjacency, it suffices to represent this notion in terms of the functions  $north : V \rightarrow V$  and  $east : V \rightarrow V$ , which return the variable corresponding to the scene that is north (east) of the designated variable.

The set of solutions can be ordered in terms of the objective of maximizing individual scene quality while maximizing phase difference between bases and fills and minimizing the temporal differences between (the bases of) adjacent images. Given an arbitrary solution  $s$ , its score is the value of the following weighted summation:

$$\begin{aligned}
f(s) = \Sigma_{i,j} & \\
& w_1 * ndvi(b_s(v_{i,j})) \\
& + w_2 * acca(b_s(v_{i,j})) \\
& + w_3 * ndvi(f_s(v_{i,j})) \\
& + w_4 * acca(f_s(v_{i,j})) \\
& + w_5 * date\Delta(b_s(v_{i,j}), f_s(v_{i,j})) \\
& + w_6 * cover(b_s(v_{i,j}), f_s(v_{i,j})) \\
& + w_7 * date\Delta(b_s(v_{i,j}), b_s(north(v_{i,j}))) \\
& + w_8 * date\Delta(b_s(v_{i,j}), b_s(east(v_{i,j}))) \\
& + w_9 * doy\Delta(b_s(v_{i,j}), b_s(north(v_{i,j}))) \\
& + w_{10} * doy\Delta(b_s(v_{i,j}), b_s(east(v_{i,j}))) \\
& + w_{11} * IsL5(b_s(v_{i,j})) \\
& + w_{12} * IsL7(b_s(v_{i,j}))
\end{aligned}$$

$w_1$  and  $w_2$  govern the importance of the quality of individual base images.  $w_3 - w_6$  discount the value of an image based on the quality of the fill, and on the goodness of fit between base and fill. For L5 images, where the base and fill are the same, the discount is the same as  $(w_2 + w_4)acca(b_s(v_{i,j}))$ , etc. Notice that since we assume that the temporal and spatial match between an image and itself is perfect, L5 images are not discounted on these criteria.  $w_7 - w_{10}$  deal with compatibility of bases with adjacent images (we ignore the compatibilities of fill), and  $w_{11}$  and  $w_{12}$  allow for an absolute preference for L5 or L7 images to be expressed. An optimal solution  $s^*$  to this GMG problem is one that receives the maximum score based on this function.

## Local Search Solution

Computationally, the problem solved by GMG is similar in structure to the problem of assigning frequencies to radio transmitters (Cabon *et al.* 1999), and other generalizations of the map coloring problem. There are two complete general constraint-based methods of solving such problems: through search, as with Branch and Bound algorithms; and through variable elimination, e.g. using Bucket Elimination (Dechter 2003). The worst case time and space complexity of the latter is tightly bounded by a parameter of the problem called the *induced width*, which arises out of an ordering of the variables. Specifically, complexity of Bucket Elimination is  $O(n * d^{w+1})$ , where  $n$  is the number of variables,  $d$  is the size of the largest domain, and  $w$  is the induced width. In practice, the primary drawback in performance is space; only problems with small induced width can be solved.

Given a set of variables and associated constraints, finding the ordering of the variables with a *minimum induced width* is an NP-hard problem. Although to our knowledge no proof exists, it appears that the induced width of a constraint graph arranged as a square grid of size  $n \times n$  is  $n$ . This linear growth rate imposes a practical limitation on the size of problems solved in a reasonable time by Bucket Elimination to roughly  $n = 30$ , too small for the GMG problem (Larrosa 2007). Hybrid approaches that combine BE with Branch and Bound have demonstrated an improvement in performance over pure BE for problems with a grid structure (Larrosa, Morancho, & Niso 2005). Although these results justify the future application of these methods to the GMG problem, in this effort we did not attempt to solve the problem using a complete method, but rather chose local search.

Reasons for adopting a local search method to solving constraint optimization problems are well documented: they include

1. *Anytime performance*: On average, local search behaves well in practice, yielding low-order polynomial running times (Aarts & Lenstra 1997). Because the criteria space is high-dimensional, it is difficult *a priori* to quantitatively characterize globally preferred solutions. Consequently, our customers were interested in a system that could examine large parts of the search space quickly to determine weight settings that produced adequate results.
2. *Flexibility and ease of implementation*: Our customers required us to build, and demonstrate the advantages of, automated solutions in a short period of time (2 months). Local search can be easily implemented.
3. *Ability to solve large problems*: As optimization problems go, the GMG-COP can be considered large. Local search has been shown to be effective on large problems.

Local search defines a class of approximation algorithms that can find near-optimal solutions within reasonable running times. Given the pair  $(S, f)$ , where  $S$  is the set of solutions and  $f$  is the objective function, let  $S^*$  be the set of best solutions (i.e., the ones with the highest score according to  $f$ ), and  $f^*$  be the best score. Members of  $S^*$  are called *global optima*. Local search iteratively searches through the set  $S$  to find a *local optimal* solution, a solution for which

no better can be found. Local optima need not be in general members of  $S^*$ .

To conduct the search, local search relies on the notion of a *neighborhood function*,  $N : S \rightarrow S$ , a function that takes one solution (called the *current solution*) and returns a new solution that differs from the current in some small way. To be effective, a neighborhood function should be simple; it should not require a lot of time to compute. For GMG-COP, the neighborhood function randomly selects a cell and replaces the selected image with a new one. A neighborhood function is *exact* if every local optima it finds as the result of search is a global optima. The neighborhood function used to solve GMG-COP is not exact.

Designing a local search algorithm is based on deciding three components: how an initial solution, or *seed* is generated, how to select a neighboring solution, and when to terminate search. For the GMG solver described in this paper, we took a simple approach to deciding these issues, reasoning that complexity should be introduced only as needed, i.e., only as warranted by inferior performance of simpler approaches, as expressed by the customers.

First, a good design for a seed generator is one that intuitively *starts in a good location* in the search space. A good location is one that is relatively close to optimal solutions, where close is measured by the length of the path from it to an optimal solution using the neighborhood function. For the GMG-COP, we chose a seed that picks the highest individual quality image for each cell, ignoring preferences related to adjacency. This seed is easy to generate (there is no need to consider adjacency constraints) and should be a good quality solution because it favors cloud-free images with high NDVI value.

Choosing a neighboring solution requires, first, choosing which cell to change. The simplest approach is to pick the cell at random. Since local search is “memoryless”, in the sense that it does not keep track of where it’s been previously, it may not be able in general to avoid examining the same solution multiple times. To avoid this, sometimes algorithms have “taboo” lists, lists of variables recently chosen to change. Variables are put on the list after chosen and eventually taken off after some number of iterations. Variables on the list can’t be selected on a given iteration. In our implementation we applied an extreme case of “taboo” list: once a scene is selected for examination, it is immediately placed on the taboo list to allow for all other scenes to be examined in the current iteration (the ordering of scene selection is random).

Given a selected cell, there are also a number of ways to select among the set of neighboring solutions based on changes made to that cell. Some are deterministic; i.e., given the same decision to make, the algorithm will make the same choice each time. Others are non-deterministic. Algorithms such as simulated annealing and genetic algorithms are non-deterministic. Initially, we opted for a deterministic approach, of which there are two kinds: first improvement or best improvement. First improvement examines neighbors, in a local search sense, until one is found that is better than the current solution; that one becomes the new current solution. Best improvement examines all the neighbors, and

picks the one that improves upon the current solution the most. Either of these generates a *greedy* approach, one that always chooses an improving solution. A variation of best improvement is where a neighbor with the best score is chosen, even if the score is worse than the score of the current solution. This approach allows for the possibility that a globally optimal solution may not be on the “greedy path” from an initial seed solution.

Finally, choosing a termination condition requires deciding how many solutions will be generated before the algorithm halts. The simplest approach will be to define a termination condition that says *halt when you reach the first locally optimal solution or after a fixed number of solutions, MAX, have been generated*, whichever comes first. A slightly more sophisticated version of this *simple local search* is called *multi-start*: here, for some fixed number of runs, we start with different initial (seed) solutions. Such initial solutions can be fully randomly generated (our implementation), semi-randomly generated, or deterministically generated. An example of deterministically generated initial solutions employed here is to assign to each scene the best self-quality image/pair. Alternatively, the local optimum of one run of simple local search can be used as the initial solution for the next run.

Testing the GMG-COP occurred in two stages. First, we compared different variations in multi-start local search to determine the best performing algorithm. Four variations were tested, based on two variations of two criteria: the initial solution and the choice of neighbor. The initial solutions tried were a randomly generated solution and the solution consisting of the set of images that scored highest individually (i.e. with respect to cloud cover, NDVI, and base-fill quality). The choice of neighbor was either done on a “first improvement” basis, i.e., the first alternative that improved the overall score, or “best improvement” basis, i.e., of all the images, selecting the one that most improved the score. The results indicate that the best strategy for finding high quality solutions is through exploration: with a random initial solution, and a “best improvement” neighbor selection, progress was quickly made towards solutions with higher quality than those found by the other approaches. We speculate that a random seed works better than one based on individual scene quality because the latter forced the search into local optimum that was not globally optimal.

In the other stage, we were interested in the extent of the improvement offered by an automated solution over current practice, which consists of manually generating solutions. Towards this end, tests were conducted by the customers at USGS and the Landsat mission using the GeoCover-2000 (GC2K) data set. The results showed that GMG, implemented as a simple algorithm then which we later improved significantly, produced a solution that was 23% better quality than the manually generated solution, based on the objective function scores. The customers viewed this result as significant enough to warrant integration and deployment of the solver.

On the complete GLS-2005 data set, the GMG solver converges to a solution in about a minute. A more detailed discussion of experiments during GMG development is found

in (Khatib *et al.* 2007).

## Landsat Scene Selection Interface (LASSI)

As noted above, the GMG solver arrives at its solution based on input consisting of a metadata representation of an image. Metadata furnish a low-fidelity, quantified assessment of several image attributes, such as cloud contamination and vegetation maturity (NDVI). Each metadata metric is a global average assessment over the entire image area, but each metric is subject to its own systematic errors (detailed discussion of which are beyond the scope of this paper). Obviously, a GMG solution is only as good as the metadata it uses to select the corresponding image, so noisy input can translate into a less than ideal solution. To address the reality of these potential sources of sub-optimal solutions, GMG is embedded into a graphical user interface and visualization tool known as “LASSI” (Large Area Scene Selection Interface). LASSI allows users to:

- Adjust objective function criteria weights prior to launching the GMG solver.
- Visually assess mosaic thumbnail image renderings of the GMG solutions.
- Examine quality of solutions with respect to each objective function criterion.
- Remotely access and view browse imagery for each WRS cell from the USGS Landsat database.

After a solution has been produced by GMG, the user may view maps representing the quality of the solution. Each map portrays a metadata attribute in color gradients on a WRS-2 map grid. These maps enable the user to assess the quality of the solution in various dimensions. Figure 1 is one such view, in this instance, of the NDVI metric. Similar metadata maps are available for

- Sensor - Discriminates Landsat 5 TM versus Landsat 7 ETM+ images in the solution set.
- Day of Year - Relative time of year of base acquisitions.
- Year - Acquisition year.
- ACCA - Cloud assessment, measure of cloud pixels (domain 0 to 10 %).
- NDVI - NDVI, normalized with respect to peak NDVI by WRS scene.
- Raw NDVI - Non-normalized (raw) NDVI.
- Preference Year - Distinguishes whether the chosen image was acquired within the middle two years (2005 or 2006) or the “fringe years” (2004 or 2007).
- Preferred Day of Year - Depicts the seasonal temporal difference between acquisition dates of the chosen image with respect to the date of the corresponding WRS scene in the GeoCover-2000 data set. (Minimizing this time difference improves the utility of the two data sets for trending analysis.)
- Northern Neighbor - Depicts the temporal difference between neighboring “along-track” scenes (north to south). Minimizing this difference reduces potential discontinuities in the resulting end-product map.

- Eastern Neighbor - Depicts the temporal difference between neighboring "adjacent-path" scenes (east to west).

Other metadata views depict the quality of images chosen to fill gaps in the L7 ETM+ images. These include quality of the ACCA fill, the cloud assessment of gap fill images; coverage, i.e., the relative success of gap-filling (100 % coverage is desired for each base-fill image pair); NDVI difference, i.e., for each base-fill image pair, the relative difference in seasonality between these images (If the difference is too great, the pair may produce a composited image with undesirable "artifacts"); and temporal difference, i.e., for each base-fill image pair, the relative difference in acquisition dates between these images (if the difference is too great, the pair may result in artifacts similar to NDVI differences, in addition to differences in sun angle).

The next visualization layer of the user interface features a thumbnail image mosaic map. By double-clicking within any metadata map, LASSI produces a mosaic map of thumbnail browse images. From here the user may view a full-screen image browse image of any thumbnail image. This is especially useful for revealing cloud contaminated images that may not be accurately represented in the metadata ACCA attribute. This mosaic map display also enables the user to view the chosen ETM+ base and fill images for side by side comparison. A small window in this display plots the monthly NDVI of this scene with a markers showing the relative acquisition time of year of the base and fill images (where applicable). Finally, the display includes a horizontally scrolling list of all candidate images of any selected WRS grid cell. From this list, the user may manually override the original GMG selection by choosing alternate acquisitions for the base and/or the fill images. Figure 2 shows a screen capture of the thumbnail image mosaic map.

Map-building using GMG/LASSI is an iterative process. The user initializes the objective function weight parameters and then invokes GMG to produce a strawman solution. After examining the metadata maps, the user tweaks these weights if necessary to compromise in some dimensions to improve others. In the end, the user will view the solution's thumbnail mosaic map and manually fine-tune it if necessary to eliminate imagery with popcorn clouds, contrails, snow, or other contaminants that may not have been accounted for in the metadata.

For the 2005 global map construction, the project has elected to pursue the problem independently for each continental landmass. This way the GMG objective function weights may be tailored for each global region. The 2005 scene selection of North America using GMG/LASSI is complete. The selection of Africa imagery is progressing. The entire global data set for the 2005 data set is planned for completion by mid-2008.

## Current Status and Future Work

Based on a successful application of the GMG on the mid-decadal global Landsat data set, the GMG will then be used for additional data set projects as well. One such mapping application has to do with the USGS goal to create a state mosaic for all 50 states in the US. The use of GMG has the

potential of automating a large part of that effort. Due to the scanner failure on L7, GMG will greatly reduce the labor necessary to exploit the Landsat data archive; and, as the L7 mission is expected to go to 2012, the benefit of the GMG cannot be overstated, due to its value in reducing the time spent examining large numbers of candidate solutions. USGS and NASA are already planning for the next Decadal survey, GLS-2010. That plan includes leveraging on the success of LASSI and GMG.

The objective function criteria and aspects of the interface occasionally undergo refinements based on evolving customer requirements. For example, a criterion was recently added that considers whether a scene is predominantly agricultural. If so, L5 imagery is preferred because the artifacts resulting from L7 image pair compositing are more noticeable and problematic when gaps are filled through homogeneous farmland. Another recent change allows for more human intervention into the solution generated. For example, as the result of a recent update, if a user manually selects a L7 base-fill image pair, then the solver is not allowed to alter that selection, or reverse the base-fill images.

The global map generation problem provides an ideal domain for testing and evaluating constraint-based optimization solvers. Furthermore, the GMG solver is of significant potential benefit to the Earth Science research community, allowing scientists access to improved automated tools to study the Earth's changing eco-system. There are future plans to apply the approach described in this paper to generating complete moon maps using Clementine image data.

## References

- Aarts, E., and Lenstra, J. K. 1997. *Local Search in Combinatorial Optimization*. John Wiley and Sons.
- Cabon, B.; de Givry, S.; Lobjois, L.; Schiex, T.; and Warners, J. P. 1999. Radio link frequency assignment. *Constraints* 4(1):78–87.
- Dechter, R. 2003. *Constraint Processing*. Morgan Kaufmann.
- Hicks, I. V.; Koster, A.; and Kolotoglu, E. 2005. Branch and tree decomposition techniques for discrete optimization. *Tutorials in Operation Research, INFORMS* 1–29.
- Khatib, L.; Gasch, J.; Morris, R.; and Covington, S. 2007. Local Search for Optimal Global Map Generation Using Mid-Decadal Landsat Images. *AAAI Workshop on Preference Handling for Artificial Intelligence*, 2007.
- Larrosa, J., and Dechter, R. 2003. Boosting search with variable elimination in constraint optimization and constraint satisfaction problems. *Constraints* 8(3):303–326.
- Larrosa, J.; Morancho, E.; and Niso, D. 2005. On the practical use of variable elimination in constraint optimization problems: 'still-life' as a case study. *JAIR* 23:421–440.
- Larrosa, J. 2007. Javier Larrosa, personal correspondence.
- Samer, M., and Szeider, S. 2006. Constraint satisfaction with bounded treewidth revisited. In Benhamou, F., ed., *Principles and Practice of Constraint Programming - CP2006*, 499–513.

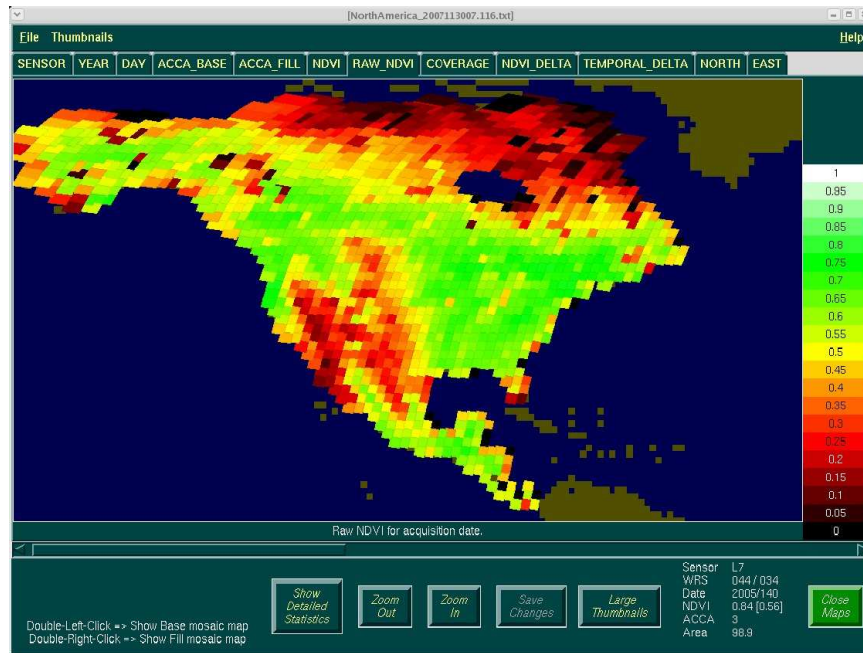


Figure 1: LASSI screen capture, showing quality of GMG-generated solution with respect to a single attribute, NDVI. Lighter color signifies better quality image with respect to attribute.

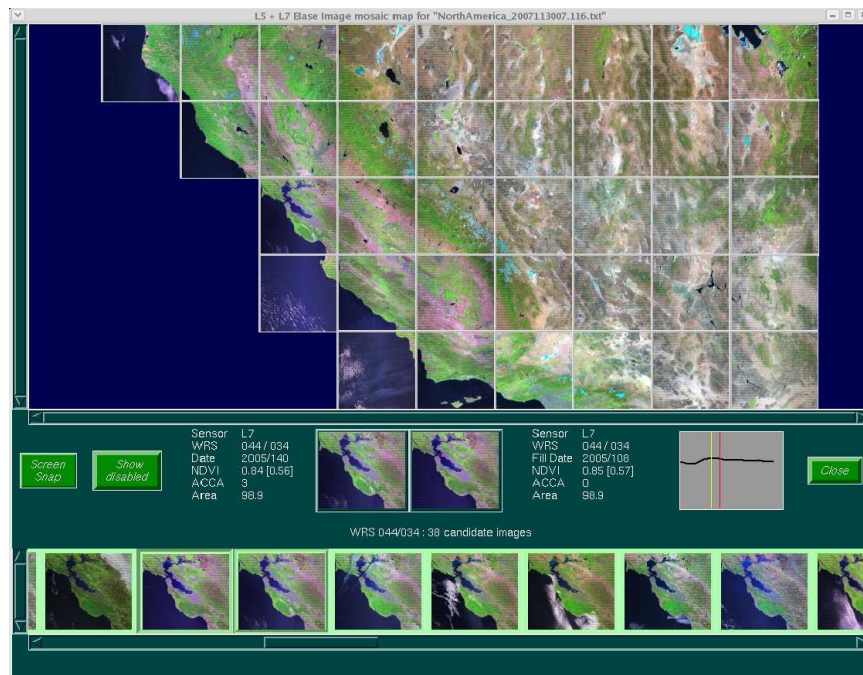


Figure 2: Screen shot of thumbnail images of GMG solution. For a selected thumbnail, the LASSI interface shows images not selected at the bottom of screen. User may override GMG selected image by choosing one of the alternatives.